

BANDWIDTH ALLOCATION SYSTEM

5 CROSS-REFERENCE TO RELATED APPLICATIONS

Not applicable.

BACKGROUND OF THE INVENTION

The present invention relates to presentation of multimedia events over a data processing network and, more particularly, to server side allocation of network bandwidth during a multimedia session comprising a layered encoded, multicast, push source and a pull source intended for multiple receivers having heterogeneous reception bandwidth.

The increased capacity and speed of the Internet and local network access have increased the practicality and popularity of delivering video conferencing, presentations, and other multimedia events over the Internet. On the other hand, the character of multimedia events and the structure of the Internet makes Internet delivery of multimedia events problematic. First, a multimedia event comprises several program elements or media, each characterized by its own qualities, delivery mechanism, and priority or importance to the presentation. Typically, a multimedia event includes audio and video, but it may also include image, text, and data elements. The various media may be classified as either a push medium where a server controls the flow and rate of data transfer or a pull medium where the data transfer rate is controlled by a client. An example of a pull medium is the HyperText Transfer Protocol (HTTP) used to transfer images, text, and, in some cases, data on the Internet. Typically, the presentation of images, text, and data is discontinuous during a multimedia session, but the priority of the data transfer is high, reliable delivery is desirable, and the bandwidth for the transmission can vary. Real-time audio and video data are typically transmitted through a multicast architecture as push media. The push media protocols

include the Real-Time Transport (RTP) protocol and the SGI (Silicon Graphics Inc.) MediaBase protocol used primarily for video-on-demand services. RTP typically operates over the User Datagram Protocol (UDP). While video has a relatively low priority, the priority of audio is high. Presentation of both audio and video is typically continuous and data rates are relatively high. While RTP/UDP provides receivers with information such as the payload type and a time stamp for reconstructing the data in real time, it does not provide any quality of service guarantees. Taken together, the elements of a multimedia event generate high data rates that can vary substantially during a session.

10 A multimedia event is also commonly an event of interest to multiple receivers and the audio and video elements are typically multicast. Multicasting enables a source to send data packets to a special multicast IP (Internet Protocol) group address that identifies a group of interested receivers. Receivers that have joined the appropriate multicast group can read the data sent to the group address. Multicasting permits a single data stream to be distributed to a number of receivers without a substantial increase in the bandwidth over that required to transmit the data to a single receiver. While the bandwidth at the source may be more than adequate for the multimedia event, the Internet is heterogeneous, comprising many connected networks, and the reception bandwidth (the bandwidth of the communication channel between a source and an individual receiver) can vary substantially between individual receivers of the multimedia event. In addition, the reception bandwidth can vary substantially over time for a particular receiver. The differences in reception bandwidth available to individual receivers means that transmission of the multimedia data at a single data rate is unlikely to be satisfactory to many of the receivers of the multicast. Allocating data rates for a plurality of media so as to preserve the integrity and character of a multimedia event and maximize the data rate at each receiver, while avoiding network congestion and data loss, is a difficult problem.

30 A solution to the problem of heterogeneous reception bandwidth for

multicast push media is the use of layered multicast combining layered data compression and heterogeneous multipoint communication. Data is encoded in several layers that can be incrementally combined when the medium is reconstructed at the receiver to provide progressively higher levels of quality of the reconstructed medium. For example, for video a base layer is encoded at a data rate that will not exceed a minimum threshold bandwidth but produce video of minimally acceptable quality at the receiver. Each additional layer improves the quality of the reconstructed video at the expense of a higher data rate. Layered multicast can be used with any push medium but is most useful for encoding video. The hierarchical dependence of the layers means that the base layer is more important than higher layers.

The multicast server simultaneously transmits the several hierarchically layered data streams to several multicast group addresses. A receiver can tune the quality of the reconstructed medium to its available reception bandwidth by adding or dropping layers from its reception. One protocol used to control layer subscription for layered multicast is the Receiver-driven Layered Multicast (RLM) protocol. Packet loss is used to determine when the receiver is over subscribed. If packet loss greater than a specified threshold is detected over a time interval known as the detection time the receiver drops the highest layer to which it is subscribed. On the other hand, in the absence of significant packet loss, a receiver will periodically experiment by joining a higher layer group until the receiver detects a significant packet loss.

While layered multicast provides a bandwidth allocation solution for push media such as video and audio, pull media, such as images, text, and data, are often important to a meaningful multimedia event and layered multicast does not provide a method of allocating bandwidth when both pull and push media are being transmitted. What is desired, therefore, is a method of allocating bandwidth during a multimedia session comprising a pull medium and a layered push medium that is multicast over a network to clients having heterogeneous reception bandwidth.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a schematic diagram of a network with multiple receivers having heterogeneous reception bandwidth.

5 FIG. 2 is a table listing certain characteristics of several media commonly included in multimedia events.

FIG. 3 is a diagram of the operation of state machine controlling reception of a layered encoded medium.

10 FIG. 4 is a graphical representation of bandwidth allocation during a multimedia event.

DETAILED DESCRIPTION OF THE INVENTION

4378944-130
T O C E T A 4 4 2 3 4 6 0

A data processing network comprises a collection of systems or nodes connected by communication links over which information can be transferred between end systems. End systems are the devices that permit users to remotely
15 access information over the network. Network applications are typically designed so that one computer or end system acts as a server managing network traffic and providing services and data to other end systems or clients attached to the network. On the other hand, network applications can provide for peer-to-peer operation where end systems may alternately assume the roles of client or server
20 from time to time. In either event, the data is transferred between end systems over a communication channel comprising links with limited data capacity or bandwidth.

Referring to FIG. 1, a complex data processing network 20, such as the Internet, comprises a number of connected networks providing a multiplicity of
25 communication links between a source 22 and a plurality of receivers, for example receivers 24, 26, 28, and 30. Each communication channel between a source and an individual receiver or group of receivers is likely to have a reception bandwidth limitation that differs from the bandwidth of other receivers connected to the network. In other words, the bandwidth of the network is heterogeneous. For
30 example, the exemplary source 22 is connected to a first router 32 by a 10 Mbps

Ethernet connection 34 (indicated by a bracket). An exemplary first receiver 24 is also connected to the router 32 by a 10 Mbps Ethernet connection 36 (indicated by a bracket) and can receive data at approximately the rate of transmission by the source 22. On the other hand, an exemplary second receiver 26 is also

5 connected to the network by a 10 Mbps Ethernet connection 38 (indicated by a bracket) but has a reception bandwidth that is limited by a 1.544 Mbps T1 link 40 (indicated by a bracket) between the first router 32 and a second router 42 in the communication channel. The reception bandwidth available to an exemplary third receiver 28 and an exemplary fourth receiver 30 are limited by the bandwidth of
10 their respective local access connections to the network, a 128 Kbps ISDN line 44 (indicated by a bracket) and a 33.6 Kbps modem 46 (indicated by a bracket).

Simultaneous transmission to all receivers at a single data rate is likely to be unsatisfactory for most of the receivers. Transmission at the lowest reception data rate would deprive receivers with higher reception bandwidth of the
15 advantages of higher data rates. In the case of video, a higher data rate reduces the distortion of the received video improving its quality. On the other hand, transmission at a higher rate would cause packet loss and loss of quality for receivers with lower reception bandwidth.

In multimedia network sessions the heterogeneity of the network is a
20 particular problem because multimedia events often require simultaneous transfer of data related to several different program elements or media having differing characteristics, delivery mechanisms, and priorities to multiple receivers connected to the network. For example, a lecture delivered as part of a distance learning activity may involve video, audio, image, data, and text elements. FIG. 2
25 lists certain exemplary characteristics 82 of several media 80 or elements that might be included in a typical multimedia network session and the communication processes used to transfer data for the medium during a session. For example, video and audio are typically transferred in a data stream "pushed" by a server while images and text are typically "pulled" by a client. The server controls the
30 flow and rate of data transfer for data in a push medium data stream. Push

protocols include Real-time Transport Protocol (RTP) for streaming media and SGI (Silicon Graphics Inc.) MediaBase for video on demand. On the other hand, the client controls the data transfer rate for a pull medium data stream by invoking a protocol such as Hypertext Transfer Protocol (HTTP) which is used for

5 transferring images, text, and, in some cases, data.

In addition to differences in the identity of the entity controlling the transfer, some media may be continuously presented while other elements of the session are presented discontinuously in discrete units. For example, video and audio elements of a session must be presented continuously to provide acceptable
10 communication. On the other hand, images are discrete elements that are presented only "occasionally." Further, the integrity and character of the multimedia event requires a different priority 84 for the delivery of the various media to an end system for presentation to a user. For example, a multimedia lecture may include slides delivered as images at appropriate points in the lecture.
15 A high transmission priority to assure prompt reception of the slides is important to the effectiveness of the lecture. Multimedia sessions including audio and video typically require high data rates, but the data rates vary substantially as media are added or deleted from the session.

Layered multicast combining layered compression and transmission of a
20 push medium provides a way of addressing the inherent bandwidth problems of a heterogeneous multipoint network connection. In a layered multicast, the source signal is divided into a plurality of layers representing incrementally increasing levels of quality. Layered multicast can be applied to data of any push medium suited to encoding in hierarchical layers but is typically applied to video. The
25 greatest compression efficiency is utilized to create a base layer having a minimal level of quality and a data rate that is satisfactory to receivers having a threshold minimum reception bandwidth. Incrementally adding layers increases the bandwidth of the transmission and the quality of the multicast video.

For example, in the network 20 the source 22 may transmit video encoded
30 from a video source information 48 in four layers; a base layer (0) 50, a second

layer (1) 52, a third layer (2) 54, and a fourth layer (3) 56. Since the exemplary first receiver 24 has sufficient reception bandwidth, it can receive all four layers and display video 58 of quality equivalent to that of the video as encoded at the source 48. On the other hand, the T1 connection 40 limits the bandwidth at the second router 42 and only the first three layers (0, 1, and 2) 50, 52, and 54 are available at the second router 42. The exemplary second receiver 26 has sufficient reception bandwidth to receive all three layers 50, 52, and 54 available at the second router 42. Since less data is available at the second receiver 26 to reconstruct the video, it will somewhat distorted (lower quality) when compared to the video "as encoded" at the source. The reception bandwidth of the exemplary fourth receiver 30 is so limited by the modem connection 46 that it is able to receive only the base layer 50 for the lowest quality display 62, while the exemplary third receiver 28 can accept two layers (0 and 1) 50 and 52 for a somewhat better quality display 64. Heterogeneity is managed by locally degrading the quality of the received signal.

The Receiver-driven Layered Multicast (RLM) protocol and the Layered Video Multicast with Retransmissions (LVMR) protocol provide rate control for multicast receivers permitting a receiver to adapt its reception data rate to the local reception bandwidth as determined by the static homogeneity of the network link and dynamic variations in link bandwidth produced by communication channel congestion. During a layered multicast the source or server 22 transmits each layer of the layered signal to a separate multicast IP (Internet Protocol) group address and takes no active role in allocating bandwidth. On the other hand, by joining or leaving a multicast group and accepting or rejecting data from the multicast group address, each receiver adapts its reception to the locally available bandwidth by adding or dropping layers of the transmission. The receiver may join a new group and add a layer of the multicast if spare (additional unused) bandwidth in the network links between the receiver and the source is apparent. When the receiver senses congestion in the communication channel, it drops a layer of the transmission.

When the receiver subscribes to a layer requiring more bandwidth than is available on the links of the network path to the receiver, the communication channel will become congested causing data packets to be dropped from the data stream degrading the quality of the received data. Data packet loss can be

5 detected by the receiver and used to instigate dropping of a layer. On the other hand, available, but unused, reception bandwidth is not readily detectable. A layered multicast receiver indirectly detects spare reception bandwidth by periodically performing a "join experiment" in which the receiver adds the next

10 oversubscribed. The receiver continues accepting data packets from the added layer unless or until it determines that the packet loss is unacceptable. If packet loss becomes unacceptable either as a result of over subscription or dynamic conditions in the communication channel, the last added layer is dropped and the receiver waits for an interval before repeating the experiment. On the other hand,

15 if packet loss remains acceptable, the receiver waits for a period and then performs another "join experiment" to determine if it can add the next hierarchical layer.

Referring to FIG. 3, when operating according to the RLM protocol, each receiver utilizes a state machine 90 to control the join experimentation. The state

20 machine 90 has four states: a steady-state 92, a hysteresis state 94, a measurement state 96, and a drop state 98. State transitions by the machine 90 are the result of either packet loss or expiration of a time interval. As a result of a state transition, the receiver can take an action as indicated parenthetically in proximity of the transition. An "add" action 100 implies subscription to the next

25 higher data rate layer in the layered multicast hierarchy and a "drop" action 102 implies dropping the highest data rate layer to which the receiver is currently subscribed and resetting a "join" timer (T_j) 104 that determines the interval that will pass before the receiver will again attempt to join the dropped layer. The relax

30 action 106 resets the join timer (T_j) 104 to decrease the interval before attempting to join the next layer. The rate of packet loss (L) either due to a sudden loss of a

single packet or an estimated sustained rate of loss is compared to a threshold rate (T). Packet loss exceeding the threshold ($L > T$) 108 causes the receiver to drop the last layer. The hysteresis state 94 provides a means of absorbing transient congestion periods. This prevents a receiver in steady state 92 from
 5 reacting to join-experiments carried out by other receivers or to transient network congestion. Measurement 96 is delayed by the period of a detection timer (T_D) 104 so congestion due to transient conditions resulting from a join-experiment will have time to dissipate.

Referring to FIG. 4, during a multimedia session a server is encoding a
 10 push medium 120 (indicated by a bracket), for example video, using a layered scheme. The push medium 120 is transmitted over the network as a base layer (0) 122 and a plurality of hierarchical, enhancement layers (1,..., i,..., L-1) 124, 126, and 128, respectively. The data rates for layers of the push medium 120 are $R_0, R_1, \dots, R_i, \dots, R_{L-1}$, respectively. At any time, a receiver or
 15 client is subscribed to the lowest i layers of the transmission having a combined data rate C_i 134 (indicated by a bracket) that does not exceed the reception bandwidth for the receiver 132. Operating under the RLM protocol, the client may add a layer as a function of the success of a join experiment or drop a layer as the result of packet loss. Since the client is receiving the push medium stream only,
 20 the total reception data rate, C_i , 134 for the client equals the sum of the individual data rates for i layers being received or:

$$C_i = \sum_{k=0}^i R_k$$

When the server begins to transmit a pull medium 136, for example an image, the total data rate at the client equals the sum of the rates for the push
 25 medium layers and the pull medium. However, the total data rate received by an individual client is limited to its reception bandwidth 132.

TCP/IP is commonly used to provide reliable delivery of image, text, and

data (pull media) on the Internet. TCP/IP includes a congestion control mechanism and transmits data at the maximum data rate available. Since the receiver is receiving layered push medium data at near the reception bandwidth, the data rate of the TCP/IP transfer of the pull medium will be suppressed and the

- 5 latency in transferring the image or other pull element will be extended. The present inventor concluded that to preserve the character of the multimedia session for network clients with heterogeneous reception bandwidth, the bandwidth of the multimedia session should be reallocated during the transmission of a pull medium. In addition, the inventor concluded that the
- 10 performance of the system could be further improved by suppressing attempts by receivers to increase the number of push medium layers being received during transmission of the pull medium.

A cost function evaluates a particular solution to a problem to determine how well the solution solves the problem. The lower the value of the cost function,

15 the better the solution to the problem. In the bandwidth allocation system of the present invention, minimization of a cost function is used to optimally reallocate the data rates of layers of a variable data rate push medium during transmission of a pull medium to optimize the bandwidth utilization and maintain the relative priorities of media to the multimedia session.

- 20 Exemplary cost functions relating distortion of a variable data rate push medium and the latency of a pull medium are:

$$\text{Cost} = D * T_j + K T_j$$

$$\text{Cost} = D + K T_j$$

where:

- 25 D = distortion of the variable data rate push medium

K = a fairness criterion

T_j = latency of the pull medium

and

$$T_j = (R_{\text{pull}} * n_j)$$

where:

R_{pull} = data rate of the pull medium

n_j = number of bits of the pull medium

5

The cost function relates the goals of simultaneously minimizing the distortion of the variable rate push medium and the latency of the pull medium. The relative contribution of each of the goals is established by the fairness criterion (K) which permits weighting of the relative importance of the distortion and the latency.

- 10 Appropriate weighting may be based on media priority specific to a particular application.

An analytical solution to the cost function requires an expression relating data rate and distortion of the medium. This expression is the result of the design of a specific encoder. If such an expression is available for the specific encoder the solution may be obtained analytically. If not, a solution can be determined by applying curve fitting to the decoder output.

15

When the server begins transmitting the pull medium, the encoding rates for the push medium layers are adjusted to encode the push medium at a combined rate of:

$$20 \quad C'_i = C_i - (R_{\text{pull_opt}})_i$$

where $(R_{\text{pull_opt}})_i$ is the optimized pull medium data rate resulting from the minimization of the cost function for the particular encoder.

- 25 To maintain the relative priority of the push medium layers, the data rate for each layer of the push medium is a function of the combined data rate for all received layers and the combined data rates for all lower priority layers:

$$R'_0 = C'_0$$

$$R'_1 = C'_1 - C'_0$$

$$R'_i = C'_i - C'_{i-1}$$

$$R'_{L-1} = C'_{L-1} - C'_{L-2}$$

When the transmission of the pull medium begins 136, the push medium server switches to encode the L layers 122, 124, 126 and 128 of the push media at the reduced rates. The encoding rates of the layers are switched back to their
 5 respective rates before reduction when the time for transmission of the pull medium has expired. If all of the bits of the pull medium are transmitted in a period inversely proportional to the available client reception bandwidth for the pull media, the transmission time equals:

$$t_i = \max\left(\frac{G * n_j}{C_i - C'_i}, T_{\min}\right)$$

10 where:

n_j = the number of transmitted bits of the pull medium

G = adjustment constant

T_{\min} = a minimum time for the rate switch

The adjustment constant G permits adjustment of the transmission time interval to
 15 reflect the fact that the actual time to transmit each pull medium will vary for each client and protocol used to transmit the medium. For example, HTTP running on top of TCP/IP initially transfers data at a slow rate but quickly permits the rate to increase to the limit of the available bandwidth. Further, each client of each multicast group will have a slightly different total reception bandwidth132 and
 20 therefore require a slightly different transmission time.

In addition to the adjustment of the push medium layer encoding rates by the server, the present inventor concluded that pull medium reception could be further improved by action at the client. Operating under the RLM and LVMR protocols, the receiver or client can perform a join experiment and attempt to add
 25 a push medium layer at any time. Since the client will reach its optimal subscription level rapidly, a layer that was added shortly before the onset of the pull medium transmission is most likely the result of a join experiment that will

ultimately fail. However, the congestion produced by a failed experiment will delay the pull medium transfer. Therefore, in the method of the present invention the highest subscribed layer of the push medium is dropped when pull medium reception begins if the push layer was "recently" added or if congestion has

5 caused the state machine 90 of the receiver to be in a state other than the steady state 92 or the drop 98 state. (If the state machine is in the drop state 98 it will progress to the steady state 92.) To determine if a layer has been "recently" added when the pull medium reception begins, each client maintains a variable for a timer that is initiated when a push layer is added. When reception of the pull
10 medium begins, the receiver compares the value of the variable to a threshold to determine if the last added layer was "recently" added. After dropping the highest layer, the state machine 90 of the receiver will advance to the "drop" state 98 and then to the steady state 92 when the drop timer T_D expires. This ensures that the state machine 90 will not overreact and drop additional layers. Once in the steady
15 state 92 further join experiments are suppressed until the pull media reception time expires.

With the method of the present invention, the data rates of the layers of a push medium of a multimedia session are reduced at the initiation of transmission of a pull medium. This increases the portion of the reception bandwidth available
20 for transfer of the pull medium promoting more rapid transfer of the pull medium. Further, suppressing join experiments during pull medium transmission reduces the likelihood of communication channel congestion which otherwise would reduce the bandwidth for the pull medium and increasing the latency of the transfer.

All the references cited herein are incorporated by reference.

25 The terms and expressions that have been employed in the foregoing specification are used as terms of description and not of limitation, and there is no intention, in the use of such terms and expressions, of excluding equivalents of the features shown and described or portions thereof, it being recognized that the scope of the invention is defined and limited only by the claims that follow.